

Data Handling The Big Picture



Kevin McFarland
CDF Offline Operations
December 19, 2001

1. Offline News and Plans for Summer 2001
2. Production
3. Impact on Needs and Priorities

http://www.pas.rochester.edu/~ksmcf/dh/dh_big_picture.pdf

Offline non-DH News

- CAF review completed
 - ↪ Not surprisingly, recommended “PC pile” architecture with network served disk cache
 - ↪ Produced a very useful “needs estimate”
 - ★ Estimated disk cache (DIM) needs at ≈ 70 TB
 - ★ Want “hundreds” of GHz processors initially, scalable to “thousands”
 - ★ Multi-branch PADs *and/or* comprehensive ntuple formats seen as key
 - ↪ Concluded that expanded fcdfsgi2 is adequate for creation of secondary datasets
 - ★ Recommends increasing the level of coordination in order to improve performance
 - ↪ Frank Wuerthwein has accepted position as Deputy Offline Leader in charge of the Central Analysis Facility

Offline non-DH News (cont'd)

- Production woes continue
 - ↪ Independent of DH problems, farms are often not able to run because of code problems
 - ↪ Integration and debugging are consuming more resources than are now available and are critical path for farms operation
 - ↪ Advertised ROOT compression failed to work on farms; result is demand for increased raw data in all streams has slowed farms peak rate down by factor of two
- PAD, multi-branch progress is slow

What about Production?

- Offline management is proposing a new model of production
 - ↪ Frozen production exe soon, with admittedly inadequate silicon reconstruction
 - ↪ Reprocess primary or secondary datasets out of production on demand
 - ↪ Secondary datasets can then split on (stable) Level-3 quantities or (stable) production output quantities
 - ↪ Plan to implement synchronously with change to 8 raw and 20 primary datasets
- Hope to operate with this model through summer 2001. Reprocessing of old (early 2001) with a fall 2001 version of production is likely.
 - ↪ Meshes well with conservative schedule for Enstore transition providing more bandwidth for farms

Offline Priorities

1. Stable real-time operation of path Level-3 → Production → Secondary Dataset Skims
 - (a) Freeze production
 - (b) Suffer to operate existing DH/tape plant
(willingly pay the price of stability, however, to migrate to Enstore as soon as commissioned)
2. Increase resources (short-term)
 - (a) Go to 8 raw and 20 primary datasets
 - (b) Increase online buffer, DIM space on fcdfsgi1, 100K bandwidth
 - (c) Increase user DIM with any available disk
(however, new network-served disk model is the long-term solution)
 - (d) fcdfsgi2 upgrade
 - (e) Seek to add “move to DIM” functionality wherever possible; however, stable operations of raw data logging, farm and migration to Enstore are higher priority
3. Move through R&D to commissioning as quickly as possible with new CAF, Enstore